arXiv:1609.03754v1 [q-bio.QM] 13 Sep 2016

Efficiency of a Stochastic Search with Punctual and Costly Restarts

Kabir Husain and Sandeep Krishna
The Simons Centre for the Study of Living Machines,
National Centre for Biological Sciences (TIFR), Bellary Road, Bangalore 560 065, India
(Dated: September 14, 2016)

The mean completion time of a stochastic process may be rendered finite and minimised by a judiciously chosen restart protocol, which may either be stochastic or deterministic. Here we study analytically an arbitrary stochastic search subject to an arbitrary restart protocol, each characterised by a distribution of waiting times. By a direct enumeration of paths we construct the joint distribution of completion time and restart number, in a form amenable to analytical evaluation or quadrature; thereby we optimise the search over both time and potentially costly restart events. Analysing the effect of a punctual, i.e. almost deterministic, restart, we demonstrate that the optimal completion time always increases proportionately with the variance of the restart distribution; the constant of proportionality depends only on the search process. We go on to establish simple bounds on the optimal restart time. Our results are relevant to the analysis and rational design of efficient and optimal restart protocols.

Stochastic searches, in which a target of interest is located by a random process, are ubiquitous in both the natural [1] and the computer sciences [2]. They may be found repeatedly in biology across a range of length scales, from the reaction kinetics of proteins in complex environments to the behavioural patterns of foragers [1, 3, 4]. These situations admit a natural interpretation as first passage time (FPT) or completion time problems [5], whose efficiency and speed is biologically or algorithmically desirable. Among several optimisation strategies, recent years has seen an interest in analysing the consequences of a 'restart' mechanism on the search process, in which the system is subject to a stochastic (or deterministic) restart while searching for the target [6–20]. Remarkably, a diverging mean completion time may be rendered finite by the introduction of a restart protocol [6, 7].

This has led to the question of how to optimise a particular search by a judiciously chosen restart process. While the bulk of existing work has considered Poisson [6–12, 19], power-law distributed [13] or deterministic [14, 17] restart protocols, there exists a large, unexplored space of possible restart mechanisms [14, 15, 20]; their study by direct enumeration is daunting. Recent work has establised that deterministic restarting *globally* minimises the mean completion time [14, 20], but physical constraints on the restart process (e.g. unavoidable stochastic fluctuations in biological systems) might make implementation of this optimal protocol unviable. Therefore, it is important to understand how introducing stochasticity into the restart process translates into changes in the mean completion time. Further, while minimising the search time may sometimes be desirable, others costs - for example, constraints on the total number of required restart events - would demand a different protocol, and this too has remained unexplored.

Here, we develop a simple yet general formalism to address these questions for a wide class of search and restart problems; we do so by solving exactly for the Laplace-transformed joint distribution of completion times and restart number. Using this result we go on to study: (i) the effect of fluctuations in the restart process on the optimal completion time, (ii) the efficiency of restart protocols when each restart event invokes a cost, and (iii) simple bounds on the globally optimal restarting time.

We obtain our results by considering the three state system depicted in Figure 1(a). This representation, which encapsulates a large class of restart problems (but not all, for example [10]), was first analysed in the context of restart processes in [18–20] by way of a renewal theory framework. Here, we undertake a complementary analysis to solve for the joint distribution of completion times and restart number.

At $t = 0$ the system begins in state A and at $t = T$ first transitions into state C. We identify the tran-

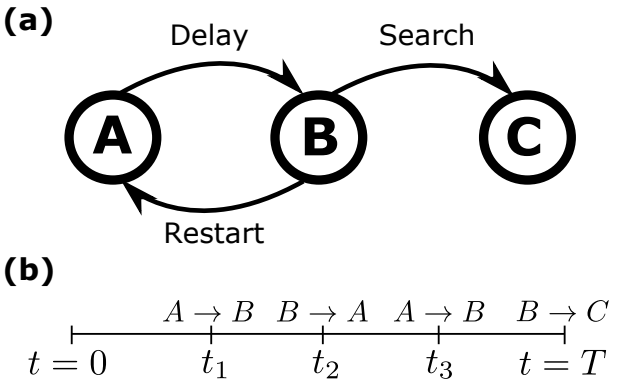

FIG. 1. Schematic of a general search process with restarts. **(a)** A 'Michaelis-Menten' representation of a stochastic search with restarts [18]. **(b)** A particular path of the system from state A at $t = 0$ to state C at $t = T$, with a single ($m = 1$) restart transition ($B \to A$).

sition $B \to C$ with that of the search (s) process, $B \to A$ with the restart (r) and $A \to B$ as a stochastic delay process (d) [19]. Each of these transitions is associated with its own waiting time distribution, $P_x(t)$ (x being s, r or d), the form of which characterises the process under study.

It is instructive to give a concrete example: in the context of the widely studied case of 1D diffusion with Poisson restarts [6, 7], $P_s(t)$ would be the Levi-Smirnov distribution: $\frac{\mathrm{e}^{-1/t}}{\sqrt{\pi}t^{3/2}}$, $P_d(t)$ would be instaneous ($= \delta(t)$) and $P_r(t)$ would be an exponential distribution with constant rate k: $k\,\exp(-kt)$. While we shall use this specific example later to illustrate our results, we will first deal with the general case, assuming only that the $P_x(t)$ are normalised (and therefore normalisable) in the usual way, but are otherwise arbitrary.

We compute the distribution of completion times $P(T)$ by an explicit summation over paths (for a recent review of similar approaches, see [21]). Consider first a particular path the system takes from A to C, wherein it undergoes the restart transition $B \to A$ m times, $m \in \{0, 1, ...\}$. Each of the $N = 2m + 2$ transitions that occur happens at a particular time t_i, where $t_{2m+2} = T$ (see Figure 1(b)). The probability density of this path may be written as:

$$P(m, T, \{t_i\}) = P_d(t_1)\left[P_r(t_2 - t_1)S_s(t_2 - t_1)\right]... \\ \times \left[P_s(T - t_{N-1})S_r(T - t_{N-1})\right], \quad (1)$$

where $S_x(t) = 1 - \int_0^t dt'\ P_x(t')$ is the *survival probability* associated with the waiting time distribution P_x [5]. Marginalising over the t_i, $1 \le i < N$, while enforcing the ordering of these intermediate times (as, in a similar context, in [11]), we obtain the joint distribution of T and m:

$$P(m, T) = \int_0^T dt_{N-1} \ldots \int_0^{t_2} dt_1\ P(m, T, \{t_i\}). \quad (2)$$

The Laplace transform of $P(m, T)$ can be compactly written as follows:

$$\mathcal{L}\left(P(m,T)\right) \equiv \int_0^\infty dT\, P(m,T)\,\mathrm{e}^{-st} = \hat{G}_f(s)\left(\hat{G}_i(s)\right)^m \quad (3)$$

where $\hat{G}_i(s) = \mathcal{L}\left[P_d(t)\right] \times \mathcal{L}\left[P_r(t)S_s(t)\right]$ and $\hat{G}_f(s) = \mathcal{L}\left[P_d(t)\right] \times \mathcal{L}\left[P_s(t)S_r(t)\right]$. Note that setting s to 0 in Eq. (3) corresponds to marginalising $P(m, T)$ over T, yielding $P(m)$. Similarly, we obtain $P(T)$ by marginalising over m:

$$\mathcal{L}\left(P(T)\right) = \hat{G}_f \sum_{m=0}^{\infty} \hat{G}_i^m = \frac{\hat{G}_f(s)}{1 - \hat{G}_i(s)}. \quad (4)$$

Eqs. (3) and (4) constitute the foundational results of our paper. Each $G_x^{(n)}$, and therefore $P(T)$, can be computed from the underlying waiting time distributions, either analytically or by quadrature. Moreover, the moments of the distribution can be obtained directly from derivatives of $\mathcal{L}\left(P(T)\right)$ at $s = 0$. When the transition from $A \to B$ is instananeous (the calculation without this assumption is similar, albeit more tedious, as is the calculation of higher moments), we may expand $\hat{G}_f$ and $\hat{G}_i$ as $\hat{G}_x(s) = G_x^{(0)} - G_x^{(1)}\,s + \frac{1}{2}G_x^{(2)}\,s^2$, where

$$G_x^{(n)} = \int dt\, t^n\, G_x(t), \quad (5)$$

and $G_i(t) = P_r(t)\ S_s(t)$ and $G_f(t) = P_s(t)\ S_r(t)$. The normalisation $\int P(T)dT = 1$ implies that $G_f^{(0)} = 1 - G_i^{(0)}$. Note that $G_i^{(0)}$ and $G_f^{(0)}$ admit a simple interpretation as the splitting probabilities for restarting (transition $B \to A$) and search completion (transition $B \to C$), respectively, from state B in Figure 1. $G_x^{(n)}/G_x^{(0)}$ is then the nth moment of the respective conditional exit time distributions [5].

Inserting (5) into (4) and expanding to first and second order in s we find:

$$\langle T \rangle = \frac{G_f^{(1)} + G_i^{(1)}}{G_f^{(0)}}, \\ \langle T^2 \rangle = \frac{2G_i^{(1)}\left(G_f^{(1)} + G_i^{(1)}\right) + G_f^{(0)}\left(G_f^{(2)} + G_i^{(2)}\right)}{(G_f^{(0)})^2} \quad (6)$$

Eqs. (4) and (6) generalise the expressions found in [19] to arbitrarily distributed restarts. We mentioned earlier that the mean completion time may be finite even if, without restarts, it diverges. Remarkably, this may be true even when the mean of the *restart* distribution also diverges. One such instance is a 1D diffusive search with disordered, random restarting; that is, restarts that occur at a constant rate k that is itself drawn (after every restart event) from, for example, an exponential distribution with mean $1/\eta$. The disorder-averaged restart distribution then decays algebraically $\propto t^{-2}$, such that the average time between restarts diverges. Nonetheless, each $G_{i/f}^{(1)}$ is finite and, from Eq. (6), so is the mean completion time [26].

A particularly interesting limit of Eq. (6) is one in which the restart occurs deterministically at a time τ, for which the mean completion time may be computed to be (as found by other means in [14, 17, 20]):

$$\langle T \rangle_\delta = \frac{\int_0^\tau dt\, S_s(t)}{1 - S_s(\tau)} \quad (7)$$

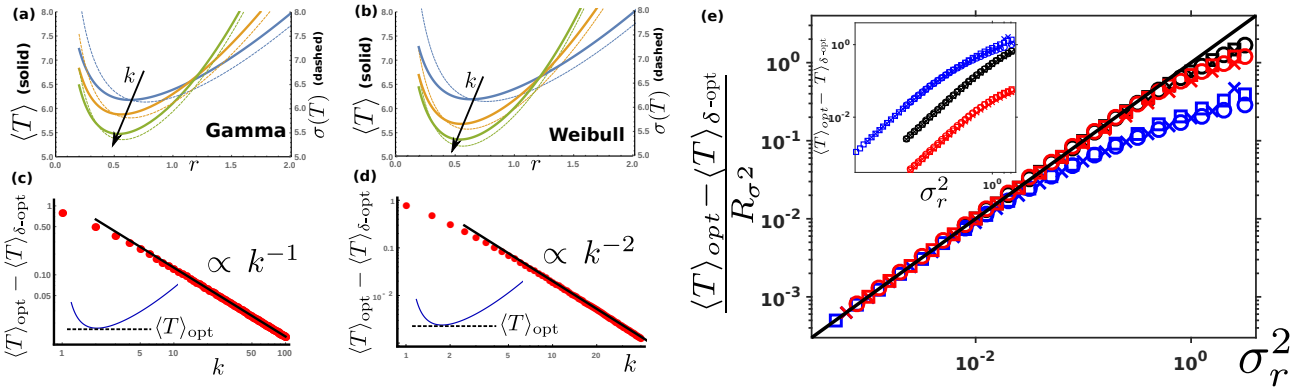

FIG. 2. The mean (solid line) and standard deviation (dashed line) of the completion time for Gamma **(a)** and Weibull **(b)** restarting with mean restart time $1/r$, with shape parameters $k = 1$ (blue), $k = 2$ (orange) and $k = 10$ (green). **(c)** and **(d)** show an algebraic decay of the optimal completion time ($\langle T\rangle_{\rm opt}(k)$), for Gamma and Weibull restarting respectively, as a function of shape parameter k. **(e)** Optimal mean completion time $\langle T\rangle_{\rm opt}$, scaled by the response factor R_{σ^2} (see Eq. (9)), as a function of the variance of the restart distribution. Solid line is a straight line with slope 1 and intercept 0. Inset shows the same without the scaling. Search distributions $P_s(t)$ shown are: Levi-Smirnov (black), log-normal (red) and Frechet (blue); for parameter values see supplementary information [26]. Restart distributions $P_r(t)$ are: Gamma (crosses), Weibull (circles) and truncated normal (squares). Data points were generated by quadrature and numerical optimisation using Eqs. (5) and (6).

This expression is finite for any τ for which $S_s(\tau) < 1$, which leads to the remarkable conclusion that a deterministic restart (optimal or not) may *always* render the mean completion time finite. Further, $\langle T\rangle_{\delta\text{-opt}} \equiv \min_\tau \langle T\rangle_\delta$ is in fact the *globally* optimal restart time, over all restart distributions [14, 20], which is achieved at the restart time $\tau_{\rm opt} = \arg\min_\tau \langle T\rangle_\delta$. How close to this optimum does a stochastic restart mechanism get? To understand the landscape of optimal completion times, we go on to ask how the introduction of fluctuations to a deterministic restart mechanism (that is, varying the shape of the restart distribution 'away' from a δ-function) affects the optimality of the completion time.

To do so in a consistent manner, we recall that several families of distributions interpolate smoothly between a δ-function (deterministic restarting) and an exponential distribution (Poisson restarting). Two illustrative examples are the Gamma family of distributions, $\frac{(kr)^k t^{k-1}}{\Gamma(k)} {\rm e}^{-krt}$, and the Weibull family, $\frac{k}{\lambda^k} t^{k-1} \exp\left(-\left(\frac{t}{\lambda}\right)^k\right)$ with $\lambda = 1/r\,\Gamma\left(1+1/k\right)$, such that both distributions have mean $1/r$. Each of these families is parameterised by a 'shape factor' k such that they are exponential for $k = 1$ and approach a δ-function as $k \to \infty$.

We illustrate the effect of the shape factor k on the mean completion time of a diffusive search in Figure 2(a-b). We observe that the optimal value of $\langle T\rangle$ decreases with k [23], approaching $\langle T\rangle_{\delta\text{-opt}}$ as a power law k^α, with numerically determined exponents $\alpha \approx -1$ for Gamma restarting and $\alpha \approx -2$ for Weibull restarting (Figure 2(c-d)).

These algebraic decays, while at first appearing to be of mysterious origin, in fact also describe the variance of a Gamma or Weibull distribution as $k \to \infty$ and r is held fixed [26]. This suggests that as the restart grows more 'punctual' (i.e. the restart distribution becomes more tightly distributed, approaching a δ-function), the optimal completion time approaches $\langle T\rangle_{\delta\text{-opt}} \propto \sigma_r^2$, the variance of $P_r(t)$. We show that this is true for the general restart problem by exploiting a central moment representation to expand a generic $P_r(t)$ around its mean τ, valid for any distribution with finite moments [24]:

$$P_r(t) = \delta(t-\tau) + \sum_{n=2}^{\infty} \frac{(-1)^n}{n!} \mu_n(\tau)\, \delta^{(n)}(t-\tau) \quad (8)$$

where $\mu_n(\tau)$ is the nth central moment of the restart distribution: $\int dt\,(t-\tau)^n P_r(t)$. In particular, $\mu_2(\tau) = \sigma_r^2$, the variance of the restart distribution. As the distribution grows more peaked, we may neglect higher moments and truncate the expansion to the term $\propto \sigma_r^2$. Inserting the expansion into Eq. (6) we find $\langle T\rangle = \langle T\rangle_\delta(\tau) + R_{\sigma^2}(\tau)\sigma_r^2 + \mathcal{O}(\mu_3(\tau))$, where R_{σ^2}, which determines the response of the completion time to fluctuations of the restart process, is:

$$R_{\sigma^2} = -\frac{\langle T\rangle_\delta}{2\left(1-S_s(\tau)\right)} \left[\frac{P_s(\tau)}{\langle T\rangle_\delta} + \left.\frac{\partial P_s}{\partial t}\right|_{t=\tau}\right] \quad (9)$$

This expression, which depends *only* on the search distribution $P_s(t)$, is valid for any punctual restart distribution – in particular, we have not yet demanded

that the mean restart time τ be optimal. If we now suppose that the optimal mean restart time is approximately constant for small σ_r^2, then the *optimal* $\langle T\rangle$ will increase linearly with σ_r^2 with slope R_{σ^2}. This is borne out by Figure 2(e), in which we observe that, as $\sigma_r^2 \to 0$, $\langle T\rangle_{\rm opt}$ approaches the respective global optimum $\langle T\rangle_{\delta\text{-opt}}$ as predicted by Eq. (9).

Inspection of Eq. (9) further reveals that fluctuations in a (possible non-optimal) restart process do not always increase the mean completion time [15, 23]. However, as we know that optimal deterministic restarts ($\tau = \tau_{\rm opt}$) are indeed optimal in the entire space of restart distributions, it must be true that $R_{\sigma^2} > 0$ for $\tau = \tau_{\rm opt}$. This implies that $\partial_t P_s(\tau_{\rm opt}) < 0$, which provides for a rather simple lower bound on $\tau_{\rm opt}$ for singly peaked search distributions $P_s(t)$ (such as those considered here): $\tau_{\rm opt}$ must lie to the right of the mode of the distribution; e.g., $\tau_{\rm opt} > \alpha^2/3$ for Levi-Smirnov (where α is the initial distance to target), $\tau_{\rm opt} > (\alpha/(1+\alpha))^{1/\alpha}$ for a one parameter Frechet distribution, $\tau_{\rm opt} > \alpha((\beta-1)/(\beta+1))^{1/k}$ for log-logistic, and $\tau_{\rm opt} > e^{-\sigma^2}$ for a log-normal distribution [25].

Up until now, we have considered an 'optimal' protocol to be one that minimises the mean completion time of the search. However, the efficacy of a search protocol may not be determined by the mean completion time but instead by other constraints; for instance, constraints on the number of restart events before the search concludes. These may be in terms of a time overhead [19], which would contribute to the delay distribution $P_d(t)$, or an energetic or financial cost. We suppose that each restart event incurs a fixed cost γ, and thus consider a *cost function* linear in m: $f(m) = \gamma\, m$. Marginalising Eq. (3) over T and then averaging: $\langle f(m)\rangle \propto \langle m\rangle = G_i^{(0)}/G_f^{(0)}$.

We may now study the cost incurred by different restart protocols. Evaluating $\langle m\rangle$ for 1D diffusion with Poisson or deterministic restarts, we find that deterministic restarts incur a lower cost for restart rates r less than ≈ 1.412, but at higher r the cost rapidly outpaces that of the stochastic restart mechanism [26]. In conjunction with Eqs. (6) and (7), we may thereby identify regimes corresponding to a tradeoff between efficiency (minimising $\langle m\rangle$) and speed (minimising $\langle T\rangle$) that depend on the nature of the restart protocol (Figure 3(a)).

Finally, we consider a case in which one might be interested in *simultaneously* minimising the number of restarts and the completion time. We must construct an appropriate cost function, $f(m,T)$ that depends on both m and completion time T, and then average it over the joint distribution $P(m,T)$, Eq. (3). We consider the form $f(m,T) = m^\beta T$, with a relative weight between time and efficiency given by the exponent β.

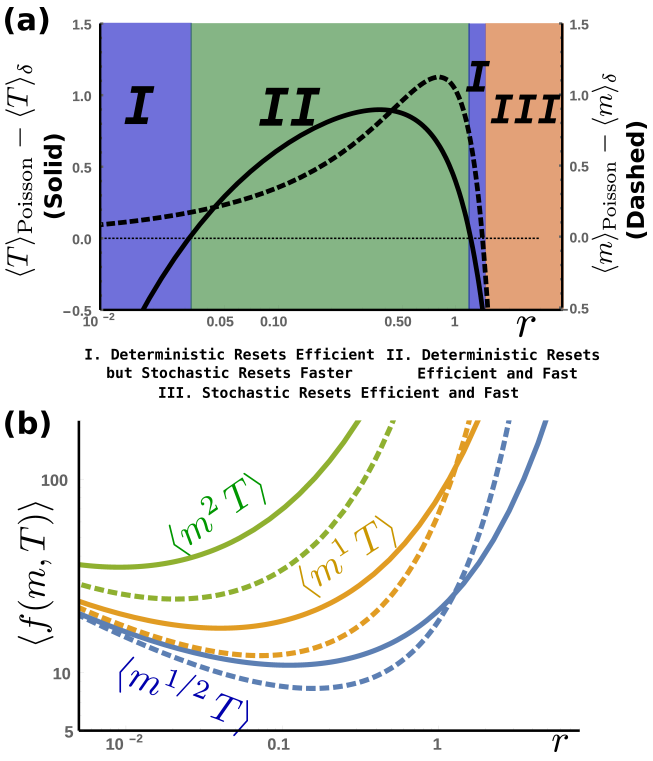

FIG. 3. **(a)** The average number of restart events for 1D diffusion with Poisson (solid) and deterministic (dashed) restarts, as a function of the inverse average restarting time ('rate') r. **(b)** The evaluated cost functions $f(m,T) = \sqrt{m}\,T$ (blue), $m\,T$ (orange) and $m^2\,T$ (green) for 1D diffusion with Poisson (solid) and deterministic (dashed) restarting with mean $1/r$.

Using Eqs. (3) and (5), we find for $\langle m^\beta T\rangle$:

$$\frac{1}{\langle m\rangle} G_f^{(0)} G_i^{(1)} \Phi\left(G_i^{(0)}, -\beta-1, 0\right) + G_f^{(1)} \Phi\left(G_i^{(0)}, -\beta, 0\right) \tag{10}$$

where $\Phi(a,b,c)$ is the Hurwitz-Lerch transcendent. This is plotted for various values of β in Figure 3(b), once again for the case of 1D diffusion with Poisson or deterministic restarts, showing that deterministic restarting continues to be the more 'optimal' protocol.

To summarize, we have introduced a simple and general framework to analyse stochastic searches with restarts, that works for arbitrary search and restart processes (and indeed, arbitrary delays $P_d(t)$). This allows us to analytically calculate (or reduce to quadratures) the moments of completion time, number of restarts, and a variety of cost functions combining these two. Thereby we are able to derive several results on the optimality of a wide range of search and restart processes. In particular, we have shown that (i) the peak of the search distribution $P_s(t)$ provides a simple lower bound for the optimal deter-

ministic restart time, $\tau_{\rm opt}$, (ii) the mean completion time of a search scales linearly with the variance of a punctual restart distribution, and (iii) deterministic restarting may invoke higher operating costs than stochastic mechanisms. Our calculations, complementary to [18–20], provide powerful tools with which to characterise the completion times of a large class of systems with restarts, and will aid the rational design of optimal, efficient restart mechanisms.

Acknowledgements We are grateful to Amit Kumar Singh, Mogens Jensen, Richard Morris and Martin Evans for their helpful comments, and B.R. Ujwal and T.E. Amit for technical advice. This work was supported by the Simons Foundation.

[1] O. Benichou et. al. Rev. Mod. Phys. **83**, 81 (2011).
[2] M. Luby, A. Sinclair and D. Zuckerman. Proceedings of the 2nd Israel Symposium on the Theory of Computing and Systems, IEEE (1993).
[3] Iddo Eliazar et al. J. Phys.: Condens. Matter **19** 065140 (2007).
[4] E. Roldan et. al. Phys. Rev. E **93**, 062411 (2016).
[5] N.G. Van Kampen. *Stochastic processes in physics and chemistry*, Elsevier (1992).
[6] M.R. Evans and S. Majumdar, Phys. Rev. Lett. **106**, 160601 (2011).
[7] M.R. Evans and S.N. Majumdar, J. Phys A. Math. Theor. **44**, 435001 (2011).
[8] M. R Evans, S.N. Majumdar and K. Mallick. J. Phys. A: Math. Theor. **46** 185001 (2013).
[9] S. Gupta, S.N. Majumdar and G. Schehr, Phys. Rev. Lett **112** 220601 (2014).
[10] S.N. Majumdar, S. Sabhapandit and G. Schehr. Phys. Rev. E, **92**, 052126 (2015).
[11] J.M. Meylahn, S. Sabhapandit, H. Touchette. Phys. Rev. E **92**, 062148 (2015).
[12] A. Pal. Phys Rev E **91**, 012113 (2015).
[13] A. Nagar and S. Gupta. Physical Review E, **93**, 060102 (2016).
[14] A. Pal et. al., J. Phys. A. Math. Theor. **49**, 225001 (2016).
[15] S. Eule and J.J. Metzger. New J. Phys., **18**, 033006 (2016).
[16] J. Fuchs, S. Goldt and U. Seifert. EPL **113**, 60009 (2016).
[17] U. Bhat, C. De Bacco and S. Redner. J. Stat. Phys. 083401 (2016).
[18] T. Rotbart, S. Reuveni and M. Urbakh, Phys. Rev. E **92**, 060101(R) (2015).
[19] S. Reuveni, Phys. Rev. Lett. **116**, 170601 (2016).
[20] A. Pal and S. Reuveni. arXiv:1607.06048 (2016).
[21] M.F. Weber and E. Frey. arXiv:1609.02849 (2016).
[22] The Gamma distribution with integer shape factor $k > 0$ and mean $1/r$ may be derived as the waiting time for k sequential Poisson events, each with rate r/k. Non-integral k follows by analytic continuation. The Weibull distribution may be similarly derived as the waiting time distribution for a time-varying rate $\propto t^{k-1}$.
[23] Note that $\langle T\rangle$ may behave non-monotonically with k when r is held fixed, as found by simulation in [15].
[24] D.T. Gillespie, Am. J. Phys. 49 (1981) 552.
[25] Associated probability densities: Levi-Smirnov: $\frac{\alpha}{\sqrt{2\pi t^3}}\,\mathrm{e}^{-\alpha^2/2t}$, Frechet: $\alpha t^{-1-\alpha}\exp(-t^\alpha)$, log-logistic: $(\beta/\alpha)(t/\alpha)^{\beta-1}/(1+(t/\alpha)^{-\beta})^2$, log-normal: $\frac{1}{\sqrt{2\pi\sigma t}}\exp\left(-(\ln t)^2/2\sigma^2\right)$.
[26] See Supplemental Material at [URL will be inserted by publisher] for some details on analytical and numerical calculations.

arXiv:1609.03754v1 [q-bio.QM] 13 Sep 2016

Efficiency of a Stochastic Search with Punctual and Costly Restarts: Supplementary Information

Kabir Husain and Sandeep Krishna
NCBS

2016

1 1D Diffusive Search

The first passage time distribution $P(t)$ for a 1D diffusive search is given by the Levi-Smirnov distribution:

$$P(t) = \frac{x_0}{\sqrt{4\pi D\, t^3}} \mathrm{e}^{-\frac{x_0^2}{4Dt}} \tag{1}$$

where x_0 is the distance to the target and D is the diffusion coefficient. We choose as the unit of time $x_0^2/4D$, with which the distribution takes the appealingly simple form $\frac{1}{\sqrt{\pi}\, t^{3/2}}\, \mathrm{e}^{-1/t}$. Note that this is a different choice than the one made by Evans and Majumdar in [1].

2 Diffusion with 'Disordered' Restarts

In the main text we construct an example in which the mean of both the search and restart distributions diverge but the mean completion time remains finite: a 1D diffusive search with Poisson restarts whose rate is drawn from an exponential distribution with mean $1/\eta$. As, to the best of our knowledge, this calculation has not been presented before, we do so here.

The restart rate k is drawn anew after each restart from the disorder distribution $P(k) = \eta\, \mathrm{e}^{-\eta\, k}$ (corresponding to a kind of 'annealed' disorder). The characteristic time-scale of restarting is then η, with $1/\eta$ being the characteristic 'rate'. Denoting an average over $P(k)$ by an overbar, we find $\overline{P}_r(t) = \eta/(\eta + t)^2$ and $\overline{S}_r(t) = \eta/(\eta + t)$. Inserting these into the definitions of $G_i(t)$ and $G_f(t)$, we find for the mean completion time:

$$\langle T \rangle = \frac{\pi\eta\, \mathrm{erfi}\left(\frac{1}{\sqrt{\eta}}\right) - 2\,{}_2F_2\left(1, 1; \frac{3}{2}, 2; \frac{1}{\eta}\right)}{1 - \frac{\sqrt{\pi} e^{\frac{1}{\eta}}\, \mathrm{erfc}\left(\frac{1}{\sqrt{\eta}}\right)}{\sqrt{\eta}}} \tag{2}$$

We plot this, as a function of η^{-1}, against the disorder-free case (simple Poisson restarts) in Figure 1(b).

3 Gamma and Weibull Distributions

In the text we analyse the properties of a diffusive search subject to restarts according to a Gamma ($P_r(t) = \frac{(kr)^k t^{k-1}}{\Gamma(k)} \mathrm{e}^{-krt}$) or Weibull ($\frac{k}{\lambda^k} t^{k-1} \exp\left(-\left(\frac{t}{\lambda}\right)^k\right)$) distribution. Note that the mean of the Gamma distribution is $1/r$ and hence varying k while keeping r fixed is easily accomplished. To facilitate comparison between the distributions, we reparameterise the Weibull distribution by the substitution $\lambda = \frac{1}{r\,\Gamma(1+1/k)}$, where the mean of the Weibull distribution is now also $1/r$.

The mean and standard deviations for each k were then calculated from the definitions of the $G_x^{(n)}$ given in the main text, either analytically (for the Gamma distribution) or by quadrature (for the Weibull distribution).

In the main text we remark that the variance of these distributions depend as power laws on the shape factor k as $k \to \infty$. We demonstrate this here. The variance of the Gamma distribution is simply $\sigma_r^2 = 1/k\, r^2$, from which the claimed dependence can be immediately seen. For the Weibull distribution, the variance is:

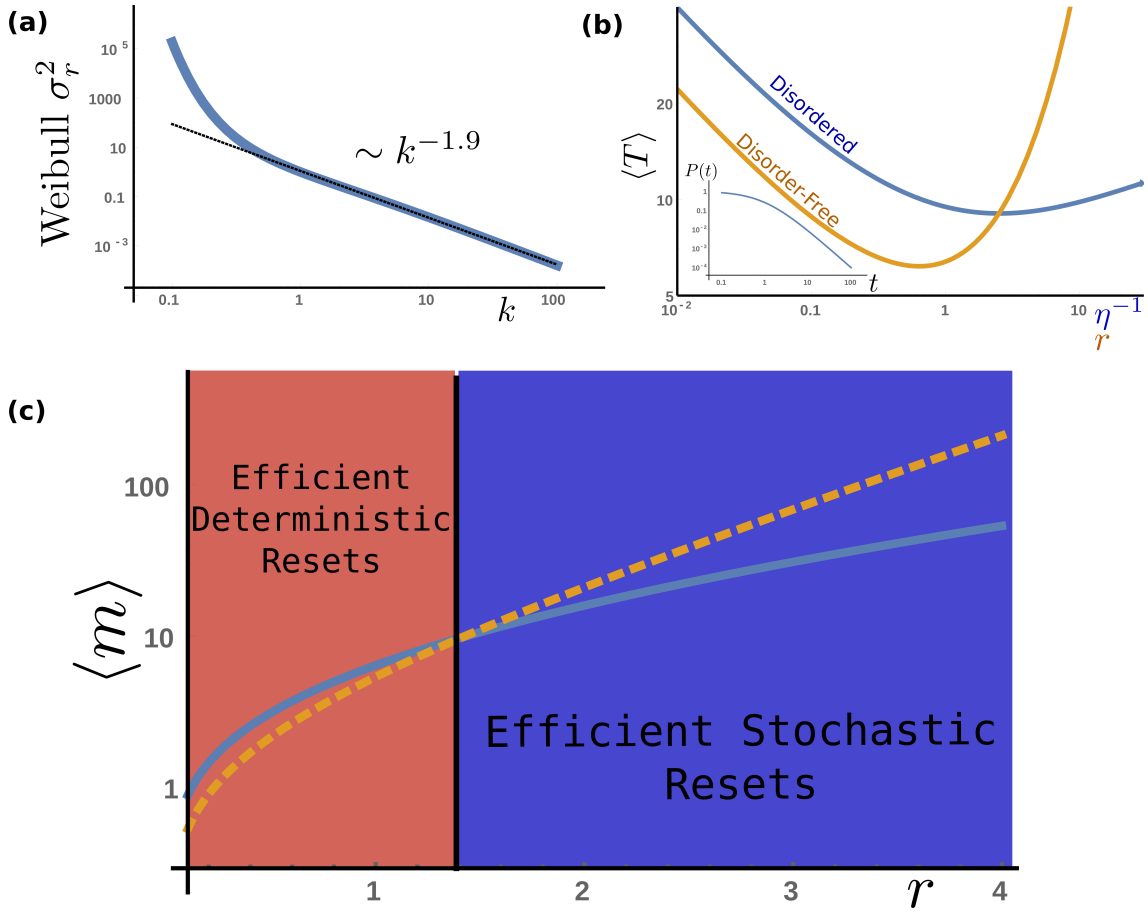

Figure 1: **(a)** Variance of a Weibull distribution σ_r^2 as as function of the shape factor k. Dotted line is $\sim k^{-1.9}$. **(b)** Average completion time $\langle T \rangle$ for Poisson (orange) restarts with rate k and exponentially disordered Poisson restarts (blue) with characteristic restart time scale η. Inset shows disorder averaged restart time distribution $\overline{P}_r(t)$ - i.e., the actual waiting time between restart events. Note the algebraic decay $\propto t^{-2}$ **(c)** Plot of $\langle m \rangle$ for a 1D diffusive search with Poisson (solid) and deterministic (dashed) restarts, as a function of the inverse mean restart time, r.

$$\sigma_r^2 = \left(\frac{1}{r\,\Gamma\left(1+1/k\right)}\right)^2 \left[\Gamma\left(1+\frac{2}{k}\right) - \left(\Gamma\left(1+\frac{1}{k}\right)\right)^2\right] \tag{3}$$

This is plotted in Figure 1(a) for a fixed r. We see that for large k the variance behaves algebraically with exponent ≈ -1.9. This was identical to the numerically determined exponent for $\langle T \rangle_{\rm opt} - \langle T \rangle_{\delta\text{-opt}}$, which we reported in the main text as ≈ -2 for simplicity.

4 Numerical Procedure for $\langle T \rangle_{\rm opt} - \langle T \rangle_{\delta\text{-opt}}$ vs. σ_r^2 Plot

For each pair of search and restart distributions, we used quadrature to find the value of $\langle T \rangle_{\rm opt}$. This was done by exploting the two-parameter nature of the restart distributions chosen. The distributions were first reparameterised by the mean $1/r$ and the variance σ_r^2. Then, for each value of σ_r^2, quadrature was used to find the mean completion time $\langle T \rangle$ for a given r – this was then numerically optimised over r to find $\langle T \rangle_{\rm opt}$. This was repeated for each value of σ_r^2, values of which were chosen so as to be uniformly distributed in log-space.

The search distributions chosen for this calculation were: Levi-Smirnov: $\frac{1}{\sqrt{\pi}\,t^{3/2}}\,\mathrm{e}^{-1/t}$, Frechet: $\frac{1}{2\,t^{3/2}}\exp\left(-1/\sqrt{t}\right)$ and log-normal: $\frac{1}{\sqrt{2\pi}\,t}\exp\left(-(\ln t)^2/2\right)$

References

[1] M.R. Evans and S. Majumdar, Phys. Rev. Lett. **106**, 160601 (2011).